

Digital Image Correlation for Performance Monitoring

Miguel Palaviccini*, Dan Turner*, and Michael Herzberg**

Abstract

Evaluating the health of a mechanism requires more than just a binary evaluation of whether an operation was completed. It requires analyzing more comprehensive, full-field data. Health monitoring is a process of nondestructively identifying characteristics that indicate the fitness of an engineered component. In order to monitor unit health in a production setting, an automated test system must be created to capture the motion of mechanism parts in a real-time and non-intrusive manner. One way to accomplish this is by using high-speed video (HSV) and Digital Image Correlation (DIC). In this approach, individual frames of the video are analyzed to track the motion of mechanism components. The derived performance metrics allow for state-of-health monitoring and improved fidelity of mechanism modeling. The results are in-situ state-of-health identification and performance prediction. This paper introduces basic concepts of this test method, and discusses two main themes: the use of laser marking to add fiducial patterns to mechanism components, and new software developed to track objects with complex shapes, even as they move behind obstructions. Finally, the implementation of these tests into an automated tester is discussed.

Introduction

The underlying premise of component health monitoring is to characterize the baseline performance of a unit, allowing for future tests to help identify any mechanical damage caused by environmental loading, as well as future performance of the component as that damage accumulates. In a production setting, the monitoring must be done in quasi real-time, with an automated system that is able to diagnose component fidelity with minimal to no human interaction. The current application focuses on using optical measurements coupled with digital image correlation (DIC) for health monitoring, as opposed to wave propagation, thermal, or other measurement variables that could also be indicative of component health.

The basic process and nomenclature of DIC is given in Ref 1. Displacements are determined by comparing an image of the deformed object to a reference image. Traditionally a square subset is used as the template to match between the reference and deformed image. In most cases, fiducials (the contrast of light and dark within the image) must be applied to the object of interest in a random fashion. The set of fiducials on any given part is often referred to as the speckle pattern.

Square subsets, while mathematically simple, can be challenging to fit onto irregularly shaped mechanism components. Conformal subsets, on the other hand, enable the entirety of the object being tracked to be included in the subset, increasing the information available for tracking and thus increasing accuracy. As a bonus, conformal subsets also enhance the tracking of objects that are partially obscured. This paper briefly discusses the advantages of the newly developed DIC code by Sandia National Laboratories, and how it is currently being used to help extract state of health information from components.

For traditional DIC applications these random fiducial patterns are often applied using a light dusting of spray paint to the product. In mechanisms work, this can be problematic. When marking a mechanism that is to be fielded, any method that adds organic material or moisture, or alters the shape and mass of the component should be avoided. Additionally, the paint method provides poor control of speckle size and placement. This paper discusses advances in laser marking techniques, along with material testing on those laser marks, as an alternate solution for adding fiducials on the components to be tracked.

* Sandia National Laboratories, Albuquerque, NM

** National Security Campus, Kansas City, MO

To summarize, traditional DIC methods rely on random contrast fiducials that are either already available on the tracked component or produced by primitive means. For mechanisms, a highly repeatable process consisting of laser marking has been developed. Furthermore, the complex geometry of the mechanisms drove the development of an algorithm that implements conformal type subsets, allowing the tracking of partially obscured objects.

DIC for Mechanisms Testing

DIC, in general, has become a popular means of determining full-field displacements from digital images. While it has also become a vital component in material characterization applications that use full-field information as part of a parameter inversion process, the main usage in mechanism testing is rigid body motion. One commercial package for DIC applications is VIC-2D, from Correlated Solutions. Analyzing motion in VIC-2D begins by defining a square subset with an odd number of pixels. The central pixel is the point of interest, and the surrounding pixels are used as information to help identify the point of interest in the subsequent video frames. This off-the-shelf DIC algorithm operates using the Lucas-Kanade method as a basis. With this, there are underlying challenges for most mechanism testing applications:

1. The image subset must have changes in intensity values in each potential axis of motion. With high-precision mechanisms components, striations from machining are often so fine as to be below the resolution of a high-speed video camera system. This is problematic, as a subset of flat color could move indeterminately within a larger area. Thus, in order to reliably track a component, intensity information must be added to the surface.
2. The second challenge is that motion of the subset must be small between frames – less than a pixel. VIC-2D interpolates the intensity information within the subset by 10x, in order to provide higher positional accuracy. The result is that movements less than 1/10th of a pixel are ideal. This dictates the minimum frame rate for video recording, and helps set camera resolution. For a regularly repeated test in a production setting, increases in frame rate and resolution increase computation time for the DIC process and storage space necessary for the resulting videos.
3. The third challenge is that a subset has a minimum size for effectiveness (about 9x9 pixels), and must be square. The minimum size requirement helps set the required camera resolution and lens, as the smallest tracked component must be capable of containing at least a 9x9 pixel square. With finer resolution, subset size can increase, but the subset must fit on the component. Thus resolution limits the total amount of information available for correlation.

In an effort to help resolve some of the above deficiencies, the tester scope was expanded to include the development of repeatable fiducials applied via laser marking, as well as the implementation of a Sandia-developed DIC software known as DICE.

Laser Marking

After an evaluation of potential methods of speckle application, laser marking was selected based on its existing use in the mechanism and low risk of creating debris or residue. Laser marking is, however, not without its concerns. The most obvious is the potential for etching or ablation of material. To address this concern, laser settings were developed in an attempt to minimize dimensional disruption of the surface. The settings must account for part material, but also for part geometry, as heat conductivity becomes more evident in parts at miniature size scales. Avoiding surface melting or ablation is a matter of keeping instantaneous power delivered to the part below a certain threshold. There are many variables involved, but a general approach is to use a defocused beam, increasing the area over which the energy is distributed.

Further discussion requires some understanding of how lasers mark stainless steel. The heating of the surface in an oxygenated environment causes growth in the chromium oxide layer. The oxide layer is

naturally transparent, but as its thickness increases, thin film effects begin to occur, and color develops. Thick enough layers absorb much of the light passing through, and appear black.

Even if no material is removed, material properties could change. The oxide layer (the bright white layer on the cross section in Figure 1) essentially behaves like a ceramic, with substantially reduced ductility compared to the surrounding metal. More critically, the laser could heat the underlying material sufficiently to alter or remove heat-treating effects in the bulk of the material. In fact, this laser marking method that avoids etching or ablation is often called 'anneal marking' or 'anneal mode'. Computer simulations predicted minimal change to part properties outside of the marked area. Metallurgical samples show an annealed zone beneath the oxide layer in the range of 1.8-2.5 μm . This is a substantial improvement from initial marking settings, which annealed up to nine times as deep. However, as a precaution, surfaces which contact other parts of the mechanism should be shielded from the laser beam to maintain material integrity. Extensive testing has been performed with marked parts without showing adverse effects.

The second concern stems from the thin film nature of the marking. Thin films are viewing-angle dependent, and can change color and contrast depending on camera and lighting angle. Laser settings that generate surface melt can also create diffraction gratings depending on line spacing. To mitigate these concerns, settings target dark gray to black marks with close line spacing and high spot overlap. Laser vendors do not generally provide all the information needed to calculate settings appropriate for a desired result, so some experimentation is necessary for optimizing settings, particularly when dealing with miniature components.

Regardless of marking method, speckle patterns must take into account the tracking algorithm and camera setup that will be used, and should be optimized to provide the best chances of success.




Figure 1: An example of laser marked surface (left) and a cross-section (right).

Pattern Design

Gradient-based DIC algorithms (like the one used in VIC-2D) perform optimally for speckles with 'soft' edges that gradually blend into the background, rather than images with sharp contrast edges which introduce high frequency content in the image. To assist with this, most gradient methods will interpolate pixel values to smooth out transitions or use an image filter. As a result, the pattern must be sized to prevent aliasing. Standard rules-of-thumb have been developed that relate the speckle to pixel size.

For X-Y motion with rotation, speckle should be at least $2 \cdot 2^{1/2}$ pixels in each dimension. Speckles larger than this size by at least one interpolation step would provide a wider peak, which is preferred. The maximum speckle size is limited by the region of interest (ROI) size. Ignoring strain, for minimum-size round speckles as above, an ROI would need at least three speckles to be fully constrained. By adding extra dimensions to speckles, it is possible to reduce the number needed to constrain an ROI. For this reason, we use ovals, or clusters of ovals, as the basic shapes for our speckle pattern. A single oval in an ROI limits it to one of two angular orientations. Local solving methods, with framerate appropriately matched to device motion, should be capable of successfully tracking such an ROI. Global methods (such

as DICE, discussed below), or test environments with insufficient framerates, may have more trouble and require more constraint.

This leads to perhaps the ultimate consideration of speckle design, which is the aperture problem. An ROI must be uniquely identifiable within the search area of the algorithm. This implies that it should be non-repetitive over that area. This was accomplished by designing a pattern that extends beyond the intended area of the region of interest. Non-repetition was verified by subjecting it to correlation tests against translated versions of itself.

For certain miniature mechanism components with complex shapes, pattern optimization was not enough to enable robust, automated tracking, and thus drove the development and incorporation of DICE for this purpose.

Benefits of DICE as the proposed tracking algorithm

The first and most substantial improvement is that subsets in DICE can be of arbitrary shape. In our application, with irregularly (non-square) regions of interest, this allows the use of more of a component's surface for matching. Conformal subsets are subsets with geometries that correspond to the outline of a particular part or region, therefore including more fiducials/speckles than traditional square subsets. An example of a conformal subset vs. a traditional square subset is shown in Figure 2 **Error! Reference source not found.** This results in an improvement in accuracy since more information is available within each conformal subset. As a bonus, this also enables the tracking of objects that are partially obscured, since even if part of a conformal subset leaves the field of view or becomes blocked by another object, there remains enough pixel information in the subset to successfully correlate.

Note that in Figure 2 it would be possible to place the markings and subsets in such a way that they fall completely within the boundary of the object of interest. If attempting to do that with a traditional square subset, the result would be either a smaller subset (yielding less information to correlate resulting in lower accuracy) or a subset that extends beyond the boundary (leading to possible de-correlation and erroneous tracking).

A second benefit to employing the algorithms in DICE is the ability to use simplex (as opposed to gradient) optimization, providing improved performance in low contrast images. However, this change in matching methods means that patterns optimized for the gradient approach of software like VIC-2D, with as many speckle edges as possible, do not perform as well in the DICE simplex method, which prefers larger areas of constant color. The lack of gradient dependence makes the simplex method useful for data sets that are impossible to analyze with the traditional Lucas-Kanade type algorithms; for example, objects without speckles, images with low contrast, and small subset size. However, the robustness comes at a cost – a substantial increase in processing time when compared to VIC-2D.

The following examples demonstrate the applicability of conformal subsets and illustrate their advantages.

Example 1: Synthetic rotation of an oddly shaped object of interest

In this example, an image data set is constructed using a synthetic sequence of images so that the exact solution is known and the error of various methods can be computed. A reference image is synthetically rotated through an angle of 10 degrees. This motion is then tracked using both square subsets and a conformal subset. The synthetic image, subsets, and results are shown in Figure 2 **Error! Reference source not found.** The magenta and green curves show the rotation error for the square subsets vs. the blue line, which shows the conformal subset results. In some cases, for example at a rotation of two degrees, the error in the square subsets is an order of magnitude higher than the conformal subset. Clearly, the accuracy of the correlation is greatly increased for the conformal subset.

Example 2: Partially obscured subsets

In this example, an obstruction is introduced into the data set that blocks a portion of the field of view, as shown in Figure 3. The obstruction is created by removing a strip of each image in the sequence. Results comparing the obstructed subset to the unobstructed results are also shown. The max difference between

the two tracked cases is 0.5 degrees, which occurs in the frame where the pawl velocity is the highest. In this case, where tracking is performed with DICE, the obstruction can be masked off and the two sections of the pawl are linked together as a single subset, leading to adequate trackable area. Note that performing this correlation using square subsets is not possible because too much of the subset becomes obstructed leading to de-correlation.




Figure 2: Synthetic image rotation used to illustrate accuracy difference between conformal and square subsets.

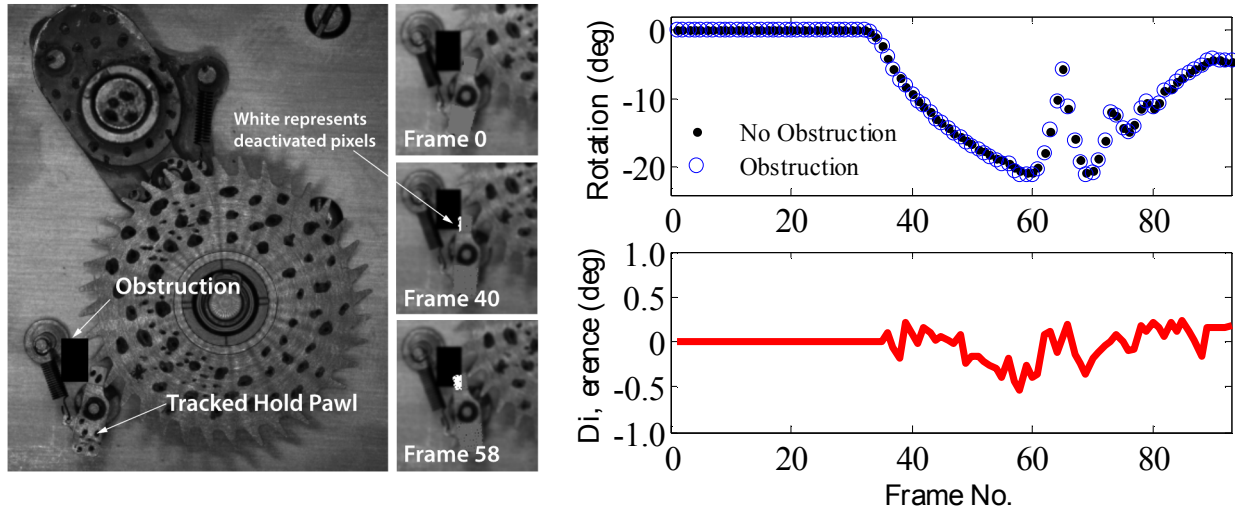


Figure 3: Rotation error computer for the synthetically generated data set using a conformal subset comparing the unobstructed and obstructed cases

Implementation into an automated tester environment

As previously mentioned, to monitor unit health in a production setting, an automated test system must be created. A system was developed to capture a video of a mechanism in operation, process the video to extract two-dimensional motion information and report results in an automated manner. This provides a statement of health prediction with a high confidence level. An implementation of the system has been in operation for several years undergoing development, and the lessons learned are currently being implemented into a qualified tester.

The tester has four main components. The first is a high speed video camera. For this application a Phantom v1210 from Vision Research along with a Canon 100 mm f/2.8 macro lens allowed for

maximizing field of view within a compact space. Note that camera and lens selection are dependent on mechanism size and motion rate. The second component is a lighting system. As with most vision solutions, the idea is to obtain as much light as possible without casting harsh shadows or creating glare off the surface of the tracked parts. The third component is a signal generator for providing the mechanism drive signals and test start trigger. Finally, a PC to allow sequencing of the test steps and for use to analyze the data as the test is run. The initial test setup was coded in LabVIEW, calling MATLAB scripts for data analysis. To speed execution and reduce maintenance, the code is now in procedural languages.

The basic order of operations is to activate the camera with a rolling frame buffer, feeding it the appropriate resolution, framerate, and exposure settings. The mechanism is then operated, with a trigger signal sent to the camera to begin saving data at the same time the mechanism receives its first drive signal. After the unit is operated, and the video has been captured, the frames are moved from the camera to the data storage folder on the computer. The computer then calls the analysis script, which feeds the video frames and predefined regions of interest to the tracking algorithm. The algorithm outputs 2-D position coordinates and angle for each ROI during each frame. These traces are then analyzed to extract time and duration of various events in the mechanism.

Conclusions and Future Work

In summary, this paper describes an improved method for analyzing mechanism performance and provides lessons learned when applying DIC to extract useful state of health metrics from components. Specifically, the paper outlines robust and repeatable methods of marking fiducials on parts of interest as well as newly developed ways of extracting information from moving parts using conformal subsets.

Laser marking allows fiducials of controlled dimensions to be added to miniature components, but presents challenges in tailoring settings to part material and shape. Laser marks are conformal for limited strain environments, and do not involve the addition of organics, moisture, or potential conductive debris required for most other conformal marking methods. Further minimizing part heating and developing patterns for low-contrast environments are ongoing topics.

As shown through examples, conformal subsets can be helpful if the tracked object is non-square by allowing more speckles to be included in the correlation. There are also a number of features for conformal subsets that are useful for trajectory tracking. For example, there are ways to enable tracking objects that cross each other's path or become partially obscured by another object. Conformal subsets can also be evolved through an image sequence to build up the intensity profile if the object is not fully visible at the start of a sequence.

With the conformal subset advantages, a price is often paid with regards to computation speed. The simplex solver method in DICe is less efficient, leading to an increase in processing time when compared to gradient-based methods. To combat this weakness, DICe has made several advances to improve parallelization and improving the performance of the simplex method. Methods of in-situ switching between fast optimization methods and more robust ones are also under development.

References

1. Sutton, M. A., Orteu, J. J., & Schreier, H. (2009). *Image Correlation for Shape, Motion and Deformation Measurements: Basic Concepts, Theory and Applications*. New York, USA: Springer.